Privacy & Security Standards Workgroup
Draft Transcript
June 10, 2011

Presentation

Judy Sparrow – Office of the National Coordinator
Good morning, everybody, and welcome to the Privacy & Security Standards Workgroup. This is a Federal Advisory Committee, so there will be opportunity at the end of the call for the public to make comment. And a reminder: Please, workgroup members, identify yourselves when speaking.

And let me do a quick roll call. Dixie Baker?

Dixie Baker – Science Applications International Corporation
I'm here.

Judy Sparrow – Office of the National Coordinator
Walter Suarez?

Walter Suarez – Kaiser Permanente
Here.

Judy Sparrow – Office of the National Coordinator
Anne Castro?

Anne Castro – BlueCross BlueShield of South Carolina
Here.

Judy Sparrow – Office of the National Coordinator
Steve Findlay? David McCallie?

David McCallie – Cerner Corporation
Here.

Judy Sparrow – Office of the National Coordinator
Wes Rishel?

Wes Rishel – Gartner, Inc.
Here.

Judy Sparrow – Office of the National Coordinator
Sharon Terry? Jeff Jonas?

Jeff Jonas – IBM
Here.

Judy Sparrow – Office of the National Coordinator
Chad Hirsch?

Chad Hirsch – Mayo Clinic
Here.

Judy Sparrow – Office of the National Coordinator
Steve Ondra? Lisa Gallagher?

Lisa Gallagher – Healthcare Information and Management Systems Society
Here.

Judy Sparrow – Office of the National Coordinator
Verne Rinker? Avinash Shanbhag?

Avinash Shanbhag – Office of the National Coordinator
Here.

Judy Sparrow – Office of the National Coordinator
Mike Davis? John Moehrke?

John Moehrke – Health Information Technology Standards Panel
Here.

Judy Sparrow – Office of the National Coordinator
Did I leave anyone off? [Pause] All right, with that, I'll turn it over to Dixie Baker and Walter Suarez.

Dixie Baker – Science Applications International Corporation
All right. Thank you, Judy, and thank you all for dialing in today to what promises to be a very interesting discussion. [Slide instruction] The agenda today and the topic of discussion really has to do with simpler, hopefully, alternatives to creating the capability to discover digital certificates as well as other provider directory-type information about provider organizations.

Before we get into the meat of our discussion, I did want to introduce a new member of our Privacy & Security Workgroup, Chad Hirsch. Chad is from Mayo Clinic and was referred to me by Chris Chute. And we're very pleased to have you on our committee, Chad, so welcome to your first meeting.

Chad Hirsch – Mayo Clinic
Thank you. Glad to be here. Look forward to working with the group.

Dixie Baker – Science Applications International Corporation
Yeah, thank you. Chad has a good, strong background in security, auditing in particular, so we're really pleased to have that kind of expertise on our workgroup.

The agenda for the day is, we'll first begin by reminding you of the context of this discussion and how we got to where we are. And then John Halamka, who—is he on the line? I didn't hear his name.

Judy Sparrow – Office of the National Coordinator
Not yet. I did send him an email, though.

Dixie Baker – Science Applications International Corporation

OK. John Halamka will introduce three concepts that we have under consideration. And one of those concepts uses the domain name service to distribute digital certificates, so we'll hear from Arien Malec on the Direct Project's experience using the domain name service for that purpose. Another of the concepts has to do with the creation of a top-level Internet domain (which is like a dot-com or a dot-edu, dot-gov, etc.) which would be the equivalent of a dot-health, if you will, and that will be an open discussion. Then David McCallie will review and discuss the concept around another of the alternatives we'll be talking about, which is the use of microformats, which is a relatively new type of delivery mechanism from Web pages. And then we'll talk about the three concepts and wrap it up and decide whether we want to make any recommendation to the Standards Committee and, if so, what that might be. [Slide instruction]

At the May HIT Standards Committee, this workgroup presented our recommendations for standards for EHR query of enterprise-level provider directories. The requirement itself for ELPDs (enterprise-level provider directories) came to us from the HIT Policy Committee, who had identified the need for a consistent approach to searching across organizations for information about that organization—directory-type information. And prime among the things that we would want them to be able to discover would be basic entity information, the information about what externally exposed services they have toward the exchange of health information, and the organization's security credentials.

So at the May meeting, we presented the first of what was likely to be a series of recommendations, and the first was the recommendations for standards for EHR query of an enterprise-level provider directory. We had considerable discussion at the committee meeting, and the full committee concluded that our recommendation was a good representation of the current state of the applicable standards in directory services, especially externally visible directory services. But they concluded that a national enterprise-level provider directory (ELPD) may not be necessary for exchange. The requirement itself, I would remind you, came to us from the Policy Committee. So the recommendation was that the Standards Committee and ONC work with the Policy Committee to refine the business requirements there. They also noted that the Direct Project, which is in—a number of people use the Direct Project protocol for exchanging [indiscernible] information—are involved in that pilot right now, and it is using the domain name service to query for digital certificates so that we're wondering how do we align the recommendation with what the Direct Project already does. [Slide instruction]

So that is a summary of what happened at the standards meeting. We did want to note that in addition to the discussion that we're having today, there are other efforts out there that are addressing provider directories. And the two that I specifically wanted to mention, just so you know that they're going on—we're not going to go into depth about them—is, the Standards and Interoperability Framework at the ONC has an initiative under way to look at standards for provider directories; and also, the grantees and the state HIE program also have a community of practice for provider directories.

So today's discussion has two objectives. One is to introduce and discuss these alternative concepts that have really been independently suggested as possible avenues for providing the kind of directory capability without implementing a full enterprise-level provider directory at the national level—and secondly to decide whether we want to recommend anything to the full Standards Committee and, if so, what that recommendation might be.

So is John on the phone now?

Judy Sparrow – Office of the National Coordinator

No, Dixie. He just sent an email. He'll be on shortly.

Dixie Baker – Science Applications International Corporation
OK. Are there any questions at this point? [Pause] OK, why don't we go ahead to the next slide?

Hopefully everybody on this call has looked at these three links we—Judy Sparrow sent them out yesterday. And they're three links to blog entries that describe the three concepts that we'll be discussing today. The first concept, which was suggested by Paul Egerman—Paul is the co-chair of the Privacy & Security Tiger Team that operates out of the Policy Committee, and his suggestion is that the health industry create a top-level domain (TLD) that is specifically for organizations that exchange health information electronically. So it would be something like mayo.health or kaiser.health. And then, once we created that TLD, we would use the same mechanism that the Direct Project uses, which is the domain name service for distributing or for lookup of digital certificates. I want to make it clear that that wouldn't be for actually issuing digital certificates; we still would need to get the digital certificates from a certificate authority. But the DNS could look up the certificate just as one uses DNS to look up the numerical address of an IP address, such as—kaiser.com translates into a numerical address. And that's described on John Halamka's blog.

The second concept that we'll be discussing is to use direct addresses as the Direct Project does today and then to use microformats that are placed on secured Web pages as a mechanism for distributing certificates and other directory information. So in other words, once you have the URL for an organization, you would go to that organization's Web page. And on that Web page would be a microformat, which is something like an address, a contact, a vCard for contact information, where you have clearly defined fields and a schema behind it. And that would be the mechanism for finding or looking up digital certificates and other information about an organization. That mechanism, which David will talk about a little later more fully, is described on Wes Rishel's blog, and that's the URL for in it on the slide there.

Then the third is really a hybrid between these two, and this concept would create the top-level domain, but it would use microformats as the mechanism for distributing directory information. The idea there is that instead of creating a full-blown directory, one could use the microformats for starting out small and then providing a richer and richer set of provider directory-type information.

Does anybody want to add anything to that just very brief top-level description of these three concepts? [Pause] OK.

Wes Rishel – Gartner, Inc.
Dixie?

Dixie Baker – Science Applications International Corporation
Yes.

Unidentified Man
I was going to suggest that Wes speak up. So Wes, speak up.

Wes Rishel – Gartner, Inc.

[Laugh] I have a question about the top-level domain. I could ask it now, and I'll let you as a chair decide whether it's appropriate to take it up now or later in the call. And that is, what is the purpose? What do we gain from having a top-level domain?

Dixie Baker – Science Applications International Corporation
The idea is that only entities that have the top-level domain would be exchanging health information through the Nationwide Health Information Network. So the vetting that would be required to be able to use that top-level domain would be a level of assurance that the entity that you're exchanging information with and that you're looking up the digital certificate is a health entity that is authorized to exchange health information on the NwHIN. That's the concept behind it.

Wes Rishel – Gartner, Inc.
OK, so then this would be an alternative to using digital certificates that were issued by a special certificate authority that would effectively be doing the same thing—that is to say, vetting the owner of the digital certificate to make sure that is a health care organization?

Dixie Baker – Science Applications International Corporation
No, it's not an alternative. It's an additional level of assurance [indiscernible] talking to a health entity. So just like we talk about defense and depth in security, you should never have one mechanism that you entrust with all your trust [laugh]. It would give us a second level of assurance, a second type of assurance, that a health entity is indeed a health entity. I mean, in truth, let's say we take the microformat approach. Anybody could put a microformat out there with a digital certificate or—

Wes Rishel – Gartner, Inc.
Yeah, I agree, and I had proposed a mechanism for assuring that it was valid that was different than the top-level domain but is used widely in the finance industry now. But let me just ask one more question: Under the top-level domain approach, we would then have two independent entities that were vetting the worthiness of an organization, one being the registrar and the other being a certificate authority. Is that right?

Dixie Baker – Science Applications International Corporation
Yes. I think that that would always have to be the case, because ICAN doesn't go out and withdraw domain names as a certificate authority can revoke a certificate if somebody proves that they aren't trustworthy of that certificate. So a CA has an ongoing kind of governance responsibility where ICAN certainly doesn't.

Wes Rishel – Gartner, Inc.
OK. So I have any number of follow-up questions, but I'm not sure if we're in this discussion right now on the agenda or you're still shaping the overall agenda.

Dixie Baker – Science Applications International Corporation
I was attempting to introduce the three concepts at a high level. And then we're going to hear more about all three of them and then have a more indepth discussion.

Wes Rishel – Gartner, Inc.
OK, so I'll reserve my other questions for the appropriate time in the agenda.

Dixie Baker – Science Applications International Corporation

Yeah, I think that probably is best, because we do have TLDs in two of the three of them, so I think that's a good idea.

David McCallie – Cerner Corporation
Dixie?

Dixie Baker – Science Applications International Corporation
Uh-huh?

David McCallie – Cerner Corporation
This is David. While we're on top-level discussion of our top-level domain (top-level orientation), I think it's important that we keep in the back of our minds the different classes of problems that these different technological solutions are being proposed for. There's a lot of overlap, but it's a slippery slope. So when we use words like "discovery," one set of ears may think of a search function, of actually finding something that they don't know based on query parameters. Another user might think of discovery in the sense of "I know what I want; I just need to get the details." And then another one may think of it as some really technical solution of a certificate from a repository somewhere. So I think we have to be careful about what problems we're solving when we talk about these various things.

And I think one clear problem is, we need to be able to get the certificate so we can use it—get the public certificate of an individual, and that's where the DNS comes in. A completely separate problem is, I need to be able to find someone's address when I have some query parameters, which is a search function, and I think that's where the notion of a Web page with embedded clues like the microformats comes into play. And then there's a third question of "What are the security implications of doing it this way?", which is where the notion of either a controlled top-level domain or EV certificates or some other approach comes into play. But be careful not to wind all those into a single fit.

Dixie Baker – Science Applications International Corporation
Yeah, I totally agree. And I think you and I discussed another way. The use of the term "discovery" really came from the Policy Committee, and they used it a bit differently from how we do from a technical sense, though. I agree with you. I think we'd probably be [indiscernible] to try to avoid using that word entirely, because it has implications beyond what we're talking here.

David McCallie – Cerner Corporation
It's not precise enough.

Dixie Baker – Science Applications International Corporation
Exactly right. I also wanted [indiscernible] that our focus here now (because I think it's really important) is on finding the digital certificate for an organization. That's the first step; that's the ELPD-type concept that we're talking about. And while, in the future, you would expect whatever this mechanism becomes to be able to accommodate lookup of individual-level directory information at the outset, we're really talking about looking for an organization's digital certificate. That's the toe in the water, if you will.

David McCallie – Cerner Corporation
Dixie, just let me push in on that. "Looking for a word" is looking for—in that you know its name and you just need its certificate or you're not even sure what its name is?

Dixie Baker – Science Applications International Corporation

I think that the assumption is that you know the name, but you don't have the digital certificate.

David McCallie – Cerner Corporation
See, I think there's another problem on the table that John's blog was addressing, which is, I'm not sure of the name. Let's say I want the cardiology department at Mayo Clinic in Rochester, and that's as much as I know, but I don't know its formal name.

Dixie Baker – Science Applications International Corporation
Yeah, that's a good point. That's another that we've discovered. A single entity like Kaiser, a very large entity, but even smaller hospitals may have multiple servers that you're looking for. So you may not know which server you're looking for, but you have a general idea of the name of the organization. So the workgroup discussed, when we were developing our recommendations, there has to be some interactive query response there.

David McCallie – Cerner Corporation
Right, and that's the point I want to make: That level of searching for or finding is different from "I know its formal name; I just need its certificate."

Dixie Baker – Science Applications International Corporation
Well, yeah, I definitely don't think we can assume that they know the specific URL for a specific server that they need to use for direct exchange, for example.

David McCallie – Cerner Corporation
Right, and that's all I wanted to make sure: We keep those two use cases separate, because they may be separate [indiscernible]. In fact, I suspect they are completely separate technical mechanisms of achieving them.

Walter Suarez – Kaiser Permanente
This is Walter. I just wanted to jump in here, because I think, as I listen to this, it gives me the impression that it's not just two use cases; it's a whole host of possible assumptions. In some cases, the assumption is, "I know the name; I don't know anything else." The other assumption might be, "I don't know the name; I just know roughly it's a Kaiser kind of entity, but I need something more to tell." So I think there are different levels of assumptions, and then there are different types of outcomes that are being looked for; at least three that I heard of are. One is, "I need to be able to find the address." "I need to be able to find the security certificate." And the third one I didn't get is, "I need to be able to know if I tracked that certificate, and I need to be able to determine whether the certificate is stressful enough." But I think we're trying to look at solutions from the three contests that, I think, address primarily one element of those, and that is the trust element of being able to find a certificate and being able to trust that the entity is a health care entity. So I'm trying to match the three things: the assumptions, the outcomes, and then the concepts.

Wes Rishel – Gartner, Inc.
This is Wes. I was agreeing with Walter 100% right up to the very last thing he said, which is that we were primarily dealing with whether the certificates were trustworthy. I believe that we're primarily dealing with all three of the types of issues that he described, because they all have to be solved and they are interlocking solutions, and that we have multiple means of solving them. And it would be incorrect to assume that there has to be a single solution that solves all three other requirements.

I'd also like to say that I believe the requirement for—let's call it finding, meaning some sort of search mechanism—for finding the address of a provider has two parts to the solution. One is the domain name of the server associate with that person. And the other is some specific name associated with the provider.

John Halamka – Harvard Medical School
Hi, this is John Halamka, very late. I couldn't get out of my doctor's office, because he was whining about meaningful use.

Wes Rishel – Gartner, Inc.
[Laugh] We're so far beyond whining about meaningful use. We're whining about much bigger things now.

David McCallie – Cerner Corporation
Or much smaller things. [Laugh]

John Halamka – Harvard Medical School
[Indiscernible] I'm leaping in here so late. Is there anything I can be helpful with in terms of background or…?

Dixie Baker – Science Applications International Corporation
Well, John, thank you for diving in. We did [indiscernible] the three concepts, and right now we're in the middle of a discussion about what actually we're searching for and the real purposes that the capability that we're discussing needs to address. And since I have the floor, I'd like to make it real clear that what we're discussing here is not the validation of certificates. The last thing that Walter mentioned was making sure that that's the right certificate. What this is is to look up a digital certificate, and any application that would use that certificate would still need to validate the certificate with the certificate authority. So I don't want it to go too far on the capabilities that we're looking for here. We really are finding the name of the server we should use and finding the digital certificate—

Wes Rishel – Gartner, Inc.
Dixie, the one phrase you're using that I'm taking exception to is "finding the name of the server we're going to use." I think a more complete formulation of that is "finding the full"—what we would call "direct site," the direct address; that is to say, it's a domain name associated with the server and some additional addressing information that allows the server. So if it's "card42" (for cardiology) "@mayo.health.us" or something, then I could look up "mayo," find some listing that Mayo had put up, and there was "cardiology lab," and I learned that it's called "card42" internally or externally. That's the user address name, because there's some guy named Card who came before them.

John Halamka – Harvard Medical School
So here's a really interesting question for the group, because the scope is really, really important to describe. When I email my friends at, say, Partners Healthcare, I happen to know their domain name is partners.org. But there is an infrastructure that can turn the fact that I know partners.org exists into an IP address and route my mail to them. Now, what if I didn't know that they were called partners.org? They could be called partnershealthcare.org or partnersaco.org. There's no mechanism for me to go searching, so to speak, on domain names to find the one that I want. And so, in this particular context, as we talk about the top-level domain, the DNS, the microformats, and all the rest, is their intent on

searching for names? Or should we assume, no, it's partners.med or partners.health or whatever the top-level domain is going to be called?

Dixie Baker – Science Applications International Corporation
That's what we were discussing: the fact that there would be—you might know Partners, but there will be multiple servers that you may be looking for. That's what Wes was just talking about—and absolutely right that you're really looking for the full domain name of the server and its full address. That's what you're looking for, but I don't think it includes—and you can argue if you will, but I don't think we're looking at the capability to go out and search for cardiologists in Redondo Beach. That's a different problem. We basically know the provider we want to find, but we may not know what exchange capabilities they have, whether they use Direct or Connect. And we may not know exactly which server we should use to exchange the information.

David McCallie – Cerner Corporation
Dixie, I think it does solve that cardiologist in Redondo Beach as a side effect. Maybe that's not a target.

Dixie Baker – Science Applications International Corporation
Right, it may be a side effect, but that's not really a part of the functional charter that we're working toward.

Walter Suarez – Kaiser Permanente
This is Walter, and the other thing is, sometimes you are not really—when I'm sending a message to an entity, I might not necessarily be directing something to a specific doctor. Or the other way to look at it is, I might be sending something to a specific doctor within a specific entity, but it is not to be sent really to the email address of that person, if you will, in the email addressing terms. It might be I'm sending a discharge summary of this patient that is with Kaiser, let's say, or with Mayo, and it's the primary care doctor—that person is Dr. Smith. But it's not a drsmith@mayo.edu or mayo.com or mayo.org or kaiser.org. It might be a site whose server, I think, has been mentioned in some of these HIE approaches that receives the external messages and then internalizes those, if you will—bring them inside and then distribute them inside to the appropriate record of the patient to the appropriate doctor. So in some cases, it's not necessarily "I need to know exactly that it's Dr. Smith and it's within a certain specific organization."

Dixie Baker – Science Applications International Corporation
In fact, at the ELPD level, we're specifically not talking about sending it to Dr. Smith. We're talking about sending it to Kaiser or Mayo or—

Wes Rishel – Gartner, Inc.
I think Walter has successfully identified four things that we're searching for. One is, we're searching for a way of addressing a person or a function. So it could be incomingdischargesummaries@mayo.edu, or it could be drmayo@mayo.edu if he's still alive. So that's two things: a service or an endpoint for distributing a message, which is a specialized kind of service. It could be we're also looking for multiple sources of assuredness that this address I'm about to use, whether it's an email address or whether it's the same combination of domain name and qualifying information construed as a URL or something else. Whatever it is, we're looking for assurance that the domain name belongs to an entity that is acceptable for exchanging health information. And the fourth thing is digital certificate retrieval, based on "I know the domain name; how do I get the retrieval?" And then off the table for today is the process for validating that that certificate is acceptable and has not been denied, has not been withdrawn.

John Moehrke – Health Information Technology Standards Panel
I guess I'd like to hand it off with that assumption of what is off the table, because if indeed we all agree that the certificate validation and revocation checking is the means by which you are assured that the certificate you got by whatever means is valid, then we don't need to add a belt and suspenders. We certainly could; i.e., the spread site has a [indiscernible]. But we do need to keep that in mind, except if we are making the assumption that certificate validation is good, we need to be in agreement that that is the presumption.

Wes Rishel – Gartner, Inc.
So Dixie addressed that as belt and suspenders, multiple lines of defense. I would lean to discuss the cost of that.

David McCallie – Cerner Corporation
And this is David. I don't—

John Moehrke – Health Information Technology Standards Panel
Yeah, I would agree. I was thinking the other thing I'd like to bring up now, since John Halamka is on now and he wasn't, Wes, when you asked the question of "What is the intended meaning or purpose or need that the top-level domain is filling?

Wes Rishel – Gartner, Inc.
What problems does the top-level domain solve?

Dixie Baker – Science Applications International Corporation
And at this point, John et al., when you do speak with—this is a public meeting, so would you please make sure you identify yourselves?

Arien Malec – Office of the National Coordinator
By the way, just for the record, I apologize. This is Arien, and I'm on as well.

Dixie Baker – Science Applications International Corporation
Hi, Arien. Thank you.

John Halamka – Harvard Medical School
Well, so this is John Halamka, and the notion was, the Standards Committee typically really likes solutions that leverage existent business processes. And so, the top-level domain was the expression of "Well, gee, there already are mechanisms out there for registrars." And so, imagine that as we start thinking about both Direct and the Nationwide Health Information Network needing these addresses that one can route to if we would establish a set of .health or .med registrars using existent business processes and put in place the controls that are already in place for other type domains, like .edu and .org, etc.—that what we would end up with is a clean list using existent business processors of endpoints that would be reasonable participants in a set of exchange activities. And it also gives a clean, technical way of having a distributed, federated directory of these folks rather than to try to shoehorn it in to .com, .org, .net, .edu.

Wes Rishel – Gartner, Inc.

So, John, you're talking now about why it would be easier to do, but you're not answering the question what value it brings.

John Halamka – Harvard Medical School
And so, what value it might bring—let's just sort of—and I'll propose this to the group—that if one challenge is, we don't want every hacker in America to be able to spew data to every single endpoint that's out there, do we, in fact, make the .med domain an effectively closed domain? And that is, the only way to use .med domain or .health domain or whatever we call it is to be a member of it so that there's some layer of spam prevention, some layer of added security one can get by creating a domain specific to this function.

Walter Suarez – Kaiser Permanente
This is Walter. I've just one question.

Doug Fridsma – Office of the National Coordinator
If I could just jump in, this is Doug Fridsma. [Indiscernible] I think one of the things that—when I think about a top-level domain—and I'm just going to give an analogy to what currently happens in the NwHIN, for example. So we have the ability to just provide certificates, and certificates are really there to provide authentication as to who that is. That certificate (that is, your digital identity) should be able to be used across different kinds of exchange. It should be able to be used to say, "I'm from Kaiser Permanente," or, "I'm from Mayo Clinic," or whatever, regardless of the method. And so, you should be able to use that in a directed exchange, and you should be able to potentially use that and say "NwHIN exchange," because it's about who you are; it's about your identity, your authentication.

The challenge that we have in the NwHIN right now is that we conflate identity and your authentication with authorization. And so, the fact that you possess a certificate that was distributed to you from the NwHIN gives you permission, then, to exchange information. If, for whatever reason, there may be a problem or you get hacked into or you're a bad actor on the NwHIN, we need to revoke your authorization. But we shouldn't necessarily revoke your authentication if that was a valid certificate [indiscernible].

So what John is articulating is that a top-level domain doesn't buy you a whole lot in terms of the issues that we're talking about with security and certificates and things like that. But it is a way, if you wanted to have a certificate for directed exchange but you didn't want to be a part of, say, this top-level domain because maybe there are some other requirements that are necessary in terms of conditions of trusted exchange or the like, you would have to get two certificates to be able to say, "Well, I'm on the NwHIN; I need a certificate for that. I want to use Direct; I need a certificate for that. I may need to do some other kind of exchange; I might need another certificate for that." In some sense, a top-level domain separates, at least from an NwHIN perspective perhaps, your authentication, associated with your certificate, and your authorization that says, "I not only have a certificate, but I've got other conditions that I've met. I get a top-level domain." Now if there's a problem, or if someone's a bad actor, you can revoke through the DNS and the domain registrars, and within 6 hours it will have promulgated throughout the world, and you are no longer authorized to have the service endpoints or to be able to function in that way. However, your digital certificate, which is your authentication and has a different mechanism to trust and obtain the trust through the route—that would still potentially be valid to say you're still OK as an entity or an organization.

So from my perspective, the top-level domain provides more value on an enforcement level than it does on the authentication/authorization piece. And so, I think when we think about the digital certificates—we think about using them in Direct or using them in NwHIN or the like—we need to think, if we conflate authorization with authentication, then we have to think about multiple certificates, because we'll need to have multiple authorizations, whereas a top-level domain actually allows you to separate those two. [Indiscernible]

John Moehrke – Health Information Technology Standards Panel
But that's not how the Internet's designed to work.

Dixie Baker – Science Applications International Corporation
Yeah.

John Moehrke – Health Information Technology Standards Panel
That's not how certificates are designed to work. That is a complete—

Dixie Baker – Science Applications International Corporation
This is John Moehrke?

John Moehrke – Health Information Technology Standards Panel
Yeah, this is John Moehrke. Turning the security world inside out and reinventing is what you're doing. You're absolutely correct that, oftentimes, certificates are used, and the fact that you have authenticated that you are the identity that the certificate represents automatically gets you authorized to do something. That's just an implementation simplicity. It's not the only way to use certificates. The fact is, as you started to say, once you have proved that somebody has authenticated who they are, you should look through your access controls engine to figure out what they are allowed to do. It's just like when you, as a user, log into a user logon session. It authorizes you to get in, but it doesn't authorize you to do everything that's in that application. So that's an implementation failure if that's how it's being done.

The other piece of this is, oftentimes, if you want to have attributes in the identity that are used to enable the access control decision (i.e., "Yes, I got the certificate, and I got it for both Direct and for exchange"), that is attributes that are put into the certificate by the certificate authority based on their provisioning of the individuals, much like when you go and get a driver's license: If you have a validation for a motorcycle, you have a validation for heavy machinery—those are all attributes that are put onto your driver's license identity. You have one identity; it has attributes that are used to prove that you are authorized to drive a motorcycle or drive heavy machinery or what have you. You don't simply say, "Well, yeah, but everybody in Wisconsin is allowed to drive heavy machinery. No one else outside Wisconsin is allowed to drive heavy machinery. Therefore, everybody has to go to Wisconsin to get a driver's license for heavy machinery." That's essentially what you've said.

Doug Fridsma – Office of the National Coordinator
No, I don't think so, because I'm trying to think about this more at an organizational level than an individual. So I'm not sure that the analogy with Wisconsin works, although I think you should continue working this through in terms of what this would mean in terms of—as somebody would need to maybe make modifications and changes, how would that roll out in terms of making those modifications to the certificates and maintaining those attributes. I think that's an important thing to explore.

John Moehrke – Health Information Technology Standards Panel
Absolutely. I'm glad we asked, because that was not the answer that we got earlier. And I think if you—to me, I don't know—Wes, I think where you were heading, this even has more to—

Wes Rishel – Gartner, Inc.
This is Wes. I've got all kinds of confusion right now about what the goal is. I think we've heard several different statements. The first one that actually made sense to me was John's discussion about spam, about being able to detect a person or a simple email filter being able to reject out of hand certain emails without having to look at them. The notion that we would do authorization backwards, as John says, puts us in the realm of invalidating a lot of existing, widely used, and widely tested open-source software and commercial software, so it is a problem in that regard. I also think that we are doing some magical thinking when we say, "Well, we know how the economics and the model work for registrars. Let's just create one—or many, perhaps, because we don't want to have a competitive lock—registrars that are specifically imbued with the additional function of vetting the identity of the applicant. I mean, I did some research about this over the weekend, and I learned how our .us top-level domain is managed (it's managed by a company under contract to the Department of Commerce) and that if you want any one of a number of domain names right now that are health related but end in a .us, you got to places like one whose company name is "$6.95 Domain Names." They actually incorporate it that way and have it be VA for a more reasonable thing, just so—it starts with an exclamation point, so the first and the last—or fatnick.com, who has similar prices (it's really his business name). And those are not people whose business model includes extensive vetting. On the other hand, certificate authorities, as I understand it, can exist across a wide range of value added and do when we understand their business model.

At any rate, whatever the business model is, the challenge we have heard is simply—and we heard this incessantly many times—is simply compiling a list of who's validated. We had the Society of State Licensing Agencies testify that their processes are locally out of date in terms of knowing who's valid, particularly who's been invalidated. Any additional information that they have beyond simply "This person has a license" is even more out of date, because physicians know they don't have to provide it in order to get their license; they just put stuff in those forms. So I think, in all of this discussion, we're in very much danger of not examining closely the fundamental thing that we're after here, which is actually having one process by which to vet an applicant for the use of these services—NwHIN or—I think the appropriate way to say it is "N-W-hin"—in various manifestations, such as Direct and Connect, much less setting up two parallel entities that are essentially doing the same function.

David McCallie – Cerner Corporation
This is David. I think that we should be really careful not to believe that our directory service, no matter how we end up formulating it, is the source of verification that somebody is a good service to do business with. I mean, the doctor may be a perfectly valid licensed physician and be a lousy physician. So we just can't get in the space of saying that if you're in the service, it means you're a good whatever. You may be awful, but you still—you may have a right to be there by whatever criteria we set, but it doesn't have anything to do with your capabilities as an entity or a provider.

Dixie Baker – Science Applications International Corporation
OK, this is Dixie. Thank you, David. Thank you, Wes. I think we should move on. I'd like to bring together those last thoughts before we move on. First of all, Wes and Walter identified the four things, I think, accurately that we're really looking for: the ability to find the person or function to direct information to or the endpoint, the service available, assurance that the domain name belongs to the right entity, and

being able to retrieve the digital certificate. I think David made a good point there that we need to avoid having our directory service as a single entity that establishes the goodness of any organization. And with that, I would really like to move on to the next slide and the next part of discussion.

The next thing on our agenda—and we've spent a lot of time here—good time, good discussion—on discussing really the benefits and challenges of TLDs, of creating a top-level directory. But I would like to go back to having Arien summarize the Direct Project experience using the domain name service as the way to retrieve digital certificates. Arien?

Arien Malec – Office of the National Coordinator
Yeah, thank you. So very briefly, as I think many people know, S/MIME requires encryption of the sent message, and it is encrypted to the encryption certificate that's used by the receiver. This gives you a number of very significant benefits, most notably that only the receiver's encryption key—private key can decrypt that message, which gives you a very high degree of assurance in the information exchange that you're participating in. The downside of that approach is that it requires the sender to know the receiver's key.

And once we decided that S/MIME was, from a standards perspective, an ideal mechanism for assuring that policy requirements that we had for direct exchange—we then had a technology challenge, which is the age-old challenge of certificate distribution. We discovered in investigation that the DNS has an existing mechanism, at least an existing specified mechanism, for placing certificates in the domain name system. In fact, the mechanism supports placing of the whole certificate or placing of a link to the certificate and a variety of mechanisms. It's called a CERT resource record, and it works reasonably well.

There are a couple of limitations to the use of the CERT record, the first being that most DNS work historically takes place over UDP, and the packets have limitation for UDP, and so it requires you to sail over to TCP in order to retrieve the certificate record. That's our experience. That is not a significant issue; that is, using TCP to retrieve domain name records is [indiscernible] and easily done.

The second issue is that the CERT record is available in most DNS software but not all. And in most particular, the Microsoft product—I'm laughing because it was actually—Shawn Ullman of Microsoft had proposed this as the enabling approach—the Microsoft DNS server does not support the CERT record, even though theoretically DNS servers are supposed to support different kinds of resource records, even ones that may not have editing capability for it. The BIND software, which is the open-source software that runs most of the DNS infrastructure, actually supports the study distribution mechanism quite well—and to get around the challenge of the Microsoft support and also to facilitate publishing.

So the other issue is that you've got a list of users and their associated [indiscernible] list of the users or organizations and their associated certificates. Getting them into your DNS software requires some stuff. It would be ideal if you could actually just serve your DNS records off of your existing mechanism, your existing storage and maintenance mechanism, for that certificate association. The team that built the reference implementations for Direct, in fact, created a lightweight DNS server that just serves the certificate off of whatever data store you have.

So the way that works is—and one of the nice attributes of DNS is that it's highly, highly, highly decentralized and federated. And there's an SRD record that—sorry, not an SRD record—an SOA record that points you to the domain name server that serves up the domain name information for a particular subdomain. So in that context, if I'm searching for direct.example.org, I look up the domain name in the

authoritative information for example.org, and it tells me that this other server actually handles direct.example.org, and that other server is in fact the server that's actually coupled with the direct reference implementation or the direct implementation.

And that works very well. It's been tested both in integration testing as well as in real-world, on-the-fly exchange. And our experience so far is that it allows for friction-free certificate distribution; and because there's a pointer to the authoritative information, it allows you to always go back to the source. So if you have, for example, an expired certificate, you can make sure that you're going back to the source of information; you have to have support caching at various levels to centralize servers. You can always invalidate your cash and invalidate the source. And then the R mechanism supports multiple listings' certificates, so I could have, for example, different trust domains for a particular user. I could have the old certificate and the newly issued certificate or the refreshed certificate multiply listed. So I've got mechanisms for multiple listings of individuals, and that ends up being quite useful.

So our experience is, there are a couple of limitations of the use of these records, most notably the reliance on TCP and the lack of availability in at least some DNS systems. We've addressed ways or mitigated those particular issues. And at least if you use BIND or one of the other DNS servers that supports CERT records or, as we would recommend, you actually couple the DNS server with the source mechanism and maintenance mechanism that you have for the association between address and certificate with the DNS server, it actually works extraordinarily well. [Indiscernible]

John Moehrke – Health Information Technology Standards Panel
This is John Moehrke. I guess I'd like to ask a question here. The question I have is, what is—as you talk about the positives and negatives—but I didn't hear you address off-the-shelf email support for this method.

Arien Malec – Office of the National Coordinator
Yeah, so that, John, is absolutely right: That is a limitation. If I'm using an off-the-shelf email client, there is good support for LDAP, usually for the internal LDAP directory that I support. There is not off-the-shelf support for DNS. If I'm using one of the reference implementations and I'm serving, for example, the DNS records straight off of my data store, it's a feasible thing, but nobody's actually tested it to do list certificates, once in LDAP and once in the DNS. So I think John's exactly right: That's another limitation of the DNS mechanism. It's not supported natively in native email clients, whereas, for example, LDAP-based certificate lookup is supported in at least some native email clients.

John Moehrke – Health Information Technology Standards Panel
As well as the reference implementation.

Arien Malec – Office of the National Coordinator
Yes. Well, the reference implementation supports DNS lookup right now for a lookup of the intra-organizational certificates, so my own certificates or the certificates that are associated with my organization. It would be a—and John and I have exchanged email about this: There is a way of listing LDAP servers in the DNS through the use of SRV records. So you point for a particular domain name (except we'll say example.org); you point "_ldap_.example.org" via an SRV record to the domain name that hosts your LDAP server. And so, that gives a mechanism for publishing LDAP information associated with a domain. And then you could, in the reference implementation, look that LDAP server up and go look at the other organization's LDAP server for that directory information. That also isn't, as far as I

know, supported in email clients, as most email clients allay to configure one or a couple of static LDAP servers that they look up.

John Moehrke – Health Information Technology Standards Panel
Yeah, the automatic discovery of an LDAP service (i.e., the SRV record) is, I don't think, all that critical, at least not in the early days, when we have a small number of directories that would be published. But I just wanted the topic to be known.

Dixie Baker – Science Applications International Corporation
OK.

Walter Suarez – Kaiser Permanente
This is Walter. I'd like to ask a couple quick questions herein. The first one is, it seems like—and it sounds like you confirmed it—you're using it as not just to transport, I guess, or to distribute the certificate but to look at certificates as well. Is that right?

Arien Malec – Office of the National Coordinator
That's correct, yes.

Walter Suarez – Kaiser Permanente
And the second question is—

John Moehrke – Health Information Technology Standards Panel
Walter, I have to ask a question: What did you mean as different between the two models?

Unidentified Man
Exactly: What's that difference?

John Moehrke – Health Information Technology Standards Panel
I didn't hear you say two different things; I heard you say the same thing [indiscernible] differently.

Walter Suarez – Kaiser Permanente
All right, so maybe the technical aspects of it are not so different. But one thing is to look at a certificate of someone, right? I'm looking for the certificate of X, whether it's an individual, or Y, whether it's an organization. Another thing is to move that certificate or retrieve that certificate. Maybe it's the same thing. [Indiscernible]

Arien Malec – Office of the National Coordinator
It ends up amounting to the same thing, with the one caveat that I could list a pointer to the certificate in my DNS record, where that certificate could be physically published someplace else.

John Moehrke – Health Information Technology Standards Panel
So that's not how we're using it today, right? We're [indiscernible].

Arien Malec – Office of the National Coordinator
That's not—it was actually the—I did add that to the specification to support both the actual DIT listing of the certificate as well as the pointer. As far as I know, nobody's using the DIT listing [indiscernible] use the pointer method.

Walter Suarez – Kaiser Permanente
[Indiscernible] The other question I have is, how much do I need to know? What are the assumptions, basically? Do I need to know the—I mean, this is one of the big questions that we have. Assumptions are what drive a lot of this. Do I know the entity? Do I know roughly the entity? Do I know the name?

Arien Malec – Office of the National Coordinator
Yeah, so the assumption—this is Arien. So there are two things that Direct uses the DNS for. One is to discover the physical machine that is used for exchange, or at least the pointer does [indiscernible] the machine or machines that are used for exchange. And that is via the DMS-MX record, which points to the mail server for the domain in question. The other is the association between the address and the certificate, and that's the topic that's under discussion. In both of those cases, the precondition, the assumption, is that I know the address. So I know I'm looking up drsmith@example.org or referrals@example.org or what have you prior to my use of the DNS in a direct scenario.

Walter Suarez – Kaiser Permanente
OK, and then my very last question is, how much is this automated? In other words, what level of human intervention, at some point, do you need? Or is this 100% automated?

Arien Malec – Office of the National Coordinator
As long as I have the address, at least as encoded in the reference implementation, all of the mechanisms are 100% automated.

Walter Suarez – Kaiser Permanente
Thank you.

Dixie Baker – Science Applications International Corporation
This is Dixie. Arien, thank you very much. One question that I have that I haven't found the answer to is that—maybe they're two questions. Number 1 is, is DMS the mechanism that has been adopted by the Direct Project? And related to that is, what is the uptake? How many people in the Direct Project actually use DNS?

Arien Malec – Office of the National Coordinator
Dixie, two very excellent questions. When we were writing the final specifications, we had a substantial debate over the degree to which the DNS was (a) a well-tested mechanism and (b) the only mechanism. And where we came down on this topic was to reflect that there's an ongoing discussion that this discussion is part of for nationwide standards for directories, that there are a number of contenders for that particular standard, and that it would be unwise for the Direct Project to say that the only way of doing certificate discovery was via the DNS. Clearly, doing a machine discovery via the DNS is exactly the right thing to do. But for the topic of certificate discovery, we thought it was appropriate to allow for experimentation, allow for people to use multiple mechanisms. At the same time (and this is the tension that we're in right now), certificate discovery is such a critical function for information exchange via Direct that unless the mechanism that is chosen is widely available, an organization that chooses something different from DNS is in effect locking its users, or at least its counterparties, into more manual methods for exchange. By the way, the other reason that we elected not to mandate DNS in the direct specification was because we explicitly wanted to support the ordinary use of S/MIME, where I physically exchange certificates prior to sending and receiving secure email, and I do that via my email client.

But many of the organizations that are deploying Direct-based services face this challenge right now of "Do I use DNS or not use DNS?" And the short-term consequences of not using DNS when everybody else is is making it difficult for users to be addressed by other organizations' users. Now, there are ways around this, so I described a mechanism for using DNS plus LDAP as a mechanism that's relatively universal. And the way that the reference implementations work, at least the way the software can work, could support multitasking or multiple ways of looking up directory information; that is, I could look up in the DNS both the CERT record and the LDAP-based SRV record and, depending on which one I find, then potentially also go out and look for LDAP. But right now, the reference implementations have great support for DNS and good support—and great support, actually, for manual methods of certificate exchange and not-so-great support for anything else.

So that's the mind, so to speak (sorry for the DNS joke), that many organizations find themselves in right now—is "DNS works, it works reasonably well, it's built into the reference implementation, and other people are adopting DNS, so why shouldn't I?" So there's no mandate of adoption, but there is a logic to organizations that are making individual choices. Now, I should also say that there are a number of organizations that believe that LDAP is the end stage solution for their provider directory. And they have, I think, a well-founded desire to distribute certificates via that LDAP mechanism, and so they're wondering, "Why should I support DNS in the short term when maybe my long-term solution is LDAP?" And the advice that I, at least, give is pretty consistently "You can support LDAP, and in fact, as I mentioned, the reference implementation, with some changes, could support LDAP reasonably transparently. But you'd need to make that investment, and other people have made the DNS investment for you. And so, it's really an economic decision about how you get started, and the easiest path right now is the DNS path." And I'm not saying that to discourage anybody from using other mechanisms, but it's really sort of an economic argument.

Wes Rishel – Gartner, Inc.
[Indiscernible] What happens if I'm an organization that has decided to use DNS and I want to send a direct communication to an organization that has decided to use LDAP?

Arien Malec – Office of the National Coordinator
Unless the reference implementation or whatever implementation that my HITSP has deployed—unless that software supports that LDAP-based lookup, then I am not going to be able to send information, unless I've out-of-band manually exchanged—

Wes Rishel – Gartner, Inc.
So "optionality" and "alternatives" are two four-letter words in this context. And [indiscernible]—just one more thing, John, before I give up: Am I interpreting you correctly to say that interorganizational LDAP is not a well-established technology yet?

Arien Malec – Office of the National Coordinator
I'm happy to be corrected here, but as far as I know, LDAP is highly broadly used and, in fact, well established within an organization. And I am not aware of significant use for interorganizational lookup. LDAP has some mechanisms for replication, but they tend to be a little touchy. And they tend to be used for—for example, one organization requires another one and needs to post a centralized directory, but we've got tangled-up instances, and so we figure out how to multiple-list and replicate and all those sorts of things.

Discovery's an issue, replication is an issue, and then—so I mentioned the DNS SRV record for LDAP. It itself isn't widely supported, although that's a little bit—the CERT record also isn't widely supported outside of Direct, but the LDAP SRV record isn't widely supported. And then the third issue for the use of DNS in an interorganizational context is that the LDAP protocol is served off of a court that is often not exposed or punched through on the firewall. So usually, I expose only my SMTP and my HTTP and my HTTPS. There's a limited set of ports that I expose my firewall; LDAP generally isn't one of them. Generally, I secure LDAP appropriately for intra-organizational lookup but not secure it and lock it down for interorganizational lookup.

Unidentified Man
Now let me just support everything Arien has just said. And all the IT organizations which I oversee—LDAP is used internally in all of them, and LDAP is used between none of them.

David McCallie – Cerner Corporation
And this is David. Conversely, DNS works as federated around the world to a remarkable degree. You can put your direct address in, register it here in Kansas City; and minutes later, someone can use it in Australia—seamless, automatic.

Unidentified Man
Dixie, has the group talked about the microformat proposal?

Dixie Baker – Science Applications International Corporation
No, and I was just about to interrupt here. We have not spoken about the microformat, and we have 7 minutes left on this call. This has been a wonderful discussion, and I thank all of you. I would propose that we schedule a follow-on talk to address the microformats, if that's OK with you and David and everybody.

David McCallie – Cerner Corporation
This is David. That's fine with me. I can do a 1-minute race through the treetops and then—

Dixie Baker – Science Applications International Corporation
OK, that sounds fine, but we do have to leave a few minutes for public comment at the end. So that's fine, David; I just didn't want to cheat you of the stage here.

David McCallie – Cerner Corporation
Yeah, and I wasn't planning to do much more than a minute or two.

Dixie Baker – Science Applications International Corporation
OK, great, all right.

David McCallie – Cerner Corporation
I'll just do a real quick run-through. This idea, I think, emerged spontaneously from a number of us: Wes and I and John Halamka and others—John Moehrke as well, recognizing that if we were to use Web pages as a way to support a query service into locating addresses, then microformats is a reasonable well-established way to embed structured information, semantic information, into the Web page such that a browser agent, a browser plug-in, or a piece of software can extract structured information out of that Web page. So, as you know, Web pages are designed for visual display; it's a layout language—

markup language. But microformats give you a way to embed the details such that you can pull them out: first name, last name, mailing address, and so forth. That's the good news.

The bad news, or at least slightly complicating news, is that there's still some controversy about the best way to do microformats. There are three separate approaches. The original's simple microformat, which is trivially easy to understand. There's one called RDFA, which uses RDF triples, which is more flexible but much more complicated and is unfortunately limited to XHTML. And then there's something called microdata, which is come out of the HTML5 WHATWG group, which is a compromise between the simple microformat model and the more complicated RDFA model. The simple model probably works find for the use case. There's something called hCard, which is a microformat mapping on top of vCard, which is an RSC standard for describing directory information. So I think in the use cases that we're proposing, the simple microformat model would probably be a good way to embed structured information into a Web page. I'll stop there.

Unidentified Man
And so, Dixie, the thing that would be most interesting in our future discussion is to compare and contrast the use of DNS versus microformats as a mechanism of certificate discovery and distribution.

Dixie Baker – Science Applications International Corporation
Right.

Wes Rishel – Gartner, Inc.
Well, I think part of what we need to do is really take Walter's four challenges that he listed and stack them up against various proposals. For example, microformats—I originally proposed that as a way to solve the problem of searching for a direct address, and it could be an additional way of finding the digital certificate if we think we need that, but I don't want it to rise or fall just as a digital certificate mechanism. It has much more potential to leverage Google and Bing and other search engines to solve a problem of getting to a national directory of self-asserted names, as opposed to ones that [indiscernible] through an as-yet-unspecified process preventing [indiscernible].

Dixie Baker – Science Applications International Corporation
I think the comparison of DNS and microformats as a way to get just certificates is not the right comparison. And I think that the four areas that Wes is referring to are the ability to find a person or a function or an endpoint, the ability to find a service, assuring that the domain name belongs to a person, and to find a digital certificate. So there are all four of those, so I do think that we need more time to discuss the microformats versus DNS for those purposes. We've talked extensively today about the top-level domain, and I personally didn't hear any compelling argument that said to me that the top-level domain is absolutely essential for either these DNS or LDAP, for that matter, or microformats to work. So if you guys agree, I think we should schedule a time just to talk about the relative merits of DNS and microformats to find the four kinds of information or perform the four functions that Walter and Wes have summarized here, if that sounds good to everybody.

John Moehrke – Health Information Technology Standards Panel
Dixie, this is John Moehrke. I certainly would support; it seems to me we have plenty to discuss. But I would also like to understand, especially with some of the leadership on the call, what the relationship to what kind of a decision we would make has to the S&I Framework sprint project on provider directory, where they are indeed looking at these additional use cases that Walter brings up and saying, "If we don't look solely at our feet but actually look a little bit out to the horizon, what would be the

choices that we would make?" And because I'm involved in both, I want to understand what the relationship of the decisions is. I'm not too interested in doing the same thing in two different places.

Arien Malec – Office of the National Coordinator
Yeah, and this is Arien, and I would second that and say that if this workgroup and the S&I Framework reach different conclusions, that would be, I think, a very bad outcome.

Dixie Baker – Science Applications International Corporation
So Arien, what would be your suggestion?

Arien Malec – Office of the National Coordinator
I think we established a process between the Standards Committee and the S&I Framework where the Standards Committee essentially hands a set of requirements or charter. And I would love to have maybe an updated charter from the Standards Committee. And I suggest that, if this group is still working through some of the essential requirements that would form part of the charter, we slow down the S&I Framework work until we have that charter. Alternatively, if we have a well-formed set of requirements, then I would suggest maybe moving some of this discussion to the S&I Framework. I guess that's the way I would frame up the decision.

Walter Suarez – Kaiser Permanente
Arien, this is Walter, and [indiscernible] need someone to help not just in this, too, but also in the policy side on this provider directory. Let me suggest one thing here, because I think, and as John pointed out, some of us are involved and we don't want to do this [indiscernible]. I suggest that the S&I Framework, which will be meeting next week face to face and continuing to pursue and [indiscernible] ideas—and many of the people are actually listening to this call—that they can take the discussion and the attention that we're paying to this microformat approach and this DNS, because they're going to be looking at the lecture. And perhaps they can bring some of that discussion back to our group in a discussion format, not in a "This is what we're recommending," but just a "This is what we have found; this is some of the things that we are looking at." And so, the S&I Framework and our workgroup maintain an ongoing back-and-forth relationship. I don't expect that those two are going to come out with different recommendations. My sense was that the S&I Framework was going to ultimately come back with a recommendation that will be brought back to us.

Arien Malec – Office of the National Coordinator
That's right. And I just don't have the experience of having, for example, the Tiger Team and the S&I Framework both explore the economic and complexity issues associated with Federal Bridge, and that ended up being really complicated. So it's better if there's a real clean separation between the work that the Standards Committee is doing and the work that the S&I Framework is doing.

Dixie Baker – Science Applications International Corporation
Well, Arien, you said that you guys established this mechanism whereby the Standards Committee hands a charter to the S&I Framework. I don't recall it ever being established, but if that's the way, I think that we need to go back to the May meeting and hand the S&I Framework the conclusion of the Standards Committee. And I can work that out with John Halamka as a co-chair of our Standards Committee. I think that's probably where we are at this point.

Arien Malec – Office of the National Coordinator
I think that works very well.

Unidentified Man
Well, but also recognize that we have talked about, as we go to formulate recommendations for the S&I Framework charter, that we can say, "Oh, we actually think there is one canonical standard that does everything. Oh, we think there's a standard that needs some work. Oh, we think you need to start from scratch." So in making those requirements, we can also make suggestions as to what we think the path forward might be.

Dixie Baker – Science Applications International Corporation
Right, and I'll remind you that some of the earlier slides—the conclusion was that the requirements need more work, so we need to be careful in how we formulate that charter.

Unidentified Man
That would be great.

Dixie Baker – Science Applications International Corporation
OK. [Indiscernible] I need to wrap this up, Walter. We're over the time limit to open it up for comments, so I'd like to do that. Judy?

Judy Sparrow – Office of the National Coordinator
Yeah. Operator, can you see if anybody wishes to make public comments?

Operator
Yes. If you are on the phone and would like to make a public comment, please press *1 at this time. If you are listening to your computer speakers, you may dial 1-877-705-2976 and press *1 to be placed in the comment queue. [Pause] We do not have any comments at this time.

Judy Sparrow – Office of the National Coordinator
OK, thank you, operator. Thank you, Dixie and everybody.

Dixie Baker – Science Applications International Corporation
And thank you all for this excellent session. Bye-bye.